

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-04-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, St Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

0254

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE May 12, 2004		3. REPORT TYPE AND DATES COVERED Final. 03/01/03 - 12/31/03	
4. TITLE AND SUBTITLE A Workshop on Information Aggregation in Decision Making				5. FUNDING NUMBERS F496200310141	
6. AUTHOR(S) Thomas S. Wallsten					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Maryland at College Park College Park, MD 20742				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				20040520 047	
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approve For Public Release: Distribution Unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) Information aggregation precedes virtually all decision making by individuals, groups, or computer algorithms. Despite the formal similarities across individual, group, and computational decision making, researchers in the three areas rarely communicate with each other. Part of the reason for the relative lack of contact between the individual-information-aggregation, group-decision-making, and optimization areas is that traditionally research in these domains has been conducted according to strict disciplinary lines. The purpose of this workshop was to foster cross-area collaboration, as well as an exchange of ideas, models, and methods. It accomplished these goals by bringing together eminent scientists whose research is focused on information aggregation in group and individual decision-making, or in associated topics in operations research, distributed detection, statistics, and formal decision science, so that the various areas of inquiry benefited from their similarities and learned from their differences. This workshop, built around both empirical and theoretical contributions, contributed toward de-compartmentalizing the research in the important areas of individual, group, and optimal decision-making.					
14. SUBJECT TERMS decision making, judgment, information aggregation				15. NUMBER OF PAGES 14	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT		

Report to AFOSR on the workshop on
"Information Aggregation in Group and Individual Decision Making"

Supported by AFOSR grant, F496200310141

This workshop took place on May 1-3, 2003 in Silver Spring, Maryland with funding from the US Air Force Office of Scientific Research, the US National Science Foundation and the *European Association of Decision Making*.

The goal of the meeting was to bring together researchers who study various aspects of information aggregation in the service of decision-making activities at different levels (individuals, small groups, large collectives), from several perspectives (normative, descriptive, prescriptive), and from different disciplinary backgrounds (psychological, statistical, computer-science, public choice). In spite of many obvious formal and substantive similarities between these lines of work, the various research perspectives have traditionally not been well integrated.

The workshop consisted of 23 presentations delivered over the course of two days. It was attended by 54 people, including the 23 presenters and 10 graduate students and post-docs. A complete list of speakers, the abstracts of their talks and Power-Point files of (almost all) the presentations can be accessed and downloaded from the workshop's web page:

<http://www.bsos.umd.edu/psyc/wallsten/workshop.htm>

The Workshop met everyone's expectations. We had hoped that bringing together leading researchers in the various areas would have a significant positive influence toward de-compartmentalizing the research. Judging by the spirited interactions, particularly among scientists from different disciplines, we believe that occurred.

DISTRIBUTION STATEMENT A

Approved for Public Release
Distribution Unlimited

Workshop on Information Aggregation in Decision Making
 Sponsored by the Department of Psychology, University of Maryland
 Administered by the University of Maryland Foundation
 May 1-3, 2003

Schedule

The Workshop will take place in the Silver Spring Hilton Hotel.
All events will take place in the Maryland Ballroom.

THURSDAY, MAY 1

5:00-8:00 PM	Reception – Maryland Ballroom (Dinner on your own)
--------------	--

FRIDAY, MAY 2

8:00 AM Continental breakfast available at the meeting room

8:30-8:45	Thomas S. Wallsten Welcome
-----------	-------------------------------

Session 1 Team Decision Making I

8:45-9:10	Kirshna Pattipati Decision Networks And Organizations
9:10-9:35	Robert D. Sorkin, S. Luan, and J. Itzkowitz Effect of majority rule and initial bias on information aggregation by groups
9:35-10:00	Verlin B. Hinsz Investigation Of A Model Of Optimal Information Integration In Groups
10:00-10:25	Barbara Mellers Implications Of A Normative Theory Of Jury Decision Making

10:25-10:55 Break

Session 2 Team Decision Making II

10:55-11:20	John Hollenbeck Asymmetric Adaptability: Dynamic Team Structures As One-Way Streets
11:20-11:45	Andrea B. Hollingshead Strategic Information Sharing In Group Decision Making
11:45-12:10	Garold Stasser Group Polarization And Collective Information Pooling: Does Pooling Promote Polarization?
12:10-12:35	John Yen CAST: Agent-based Proactive Information Exchanges for Team Decision-Makings

12:35-2:00 Buffet lunch

Session 3 Aggregation by Individuals I

2:00-2:25	Richard P. Larrick and Jack B. Soll Lay Intuitions About Combining Quantitative Judgments
2:25-2:50	David Budescu Confidence In The Aggregated Opinion Of Correlated Judges
2:50—3:15	Ilan Yaniv Receiving Advice And Making Judgments In A Social Environment: A Pragmatic Approach
3:15-3:40	Nigel Harvey Source Effects In Advice Taking

3:40-4:10 Break

Session 4 Aggregation by Individuals II

4:10-4:35	Ido Erev, Alvin E. Roth, Robert L. Slonim, and Greg Barron Combining A Theoretical Prediction With Experimental Evidence <i>Replaced By</i> Jack Soll and Richard Larrick Intuitive Strategies for Revising Opinions: Are people appropriately influenced by others?
4:35-5:00	Joseph V. Baranski Aggregating Conclusive And Inconclusive Information And The Assessment Of Threat
5:00-5:25	Cleotilde Gonzalez Learning In Dynamic Decision Making
5:25-5:50	Arndt Broeder Decisions From Memory: Information Integration Depends On The Information Format

7:00-9:30 Dinner at Mrs. K's Toll House (Directions to be provided)

SATURDAY, MAY 3

8:00 AM Continental breakfast available at the meeting room

Session 3 Collective Choice

8:30-8:55	Norman Frohlich & Joe A. Oppenheimer Justice Preferences And The Arrow Problem
8:55-9:20	Michael Regenwetter Aggregation Of Probabilistic Ordinal Preferences
9:20-9:45	Reid Hastie and Tatsuya Kameda The Robust Beauty Of The Majority Rule
9:45-10:10	Whitman Richards Graphical Models And Collective Choice

10:10-10:40 Break

Session 6 Methods of Aggregation

10:40-11:05	Dave Krantz Forecast norms, forecast errors, and the use of metacognitive judgments
11:05-11:30	Thomas S. Wallsten and David Budesu Averaging Probability Judgments
11:30-11:55	Bob Winkler Bayesian Information Aggregation: Modeling Issues And Challenges
11:55-12:20	Andrew Gelman, Robert Clemen, Roger Cooke, James Hammitt, David Krantz, And Francis Tuerlinckx Principles And Methods For Combining Expert Judgments
12:20-2:30	Buffet lunch and closing discussion

ABSTRACTS
(Alphabetically by first author)

AGGREGATING CONCLUSIVE AND INCONCLUSIVE INFORMATION AND THE ASSESSMENT OF THREAT

Joseph V. Baranski

A major concern for the naval coalition task force operating as part of Operation Apollo is to assess the potential threat of all contacts coming into and going out of the Gulf of Oman. Automated decision support systems that aid commanders in such tasks focus almost exclusively on conclusive information that favors either "hostile" or "friendly" targets. Inevitably, however, some of the information that is relevant to the judgment of threat may be inconclusive or unavailable. This presentation will provide an overview of recent work in our laboratory examining how experts and novices combine conclusive and inconclusive information when forming threat assessments using a medium fidelity simulation of a naval threat assessment operation. In a typical study, participants interrogate 10 pieces of information (e.g., speed, direction, bearing) about 'targets' in a radar space and the amount of hostile ($n(H)$), peaceful ($n(P)$), and 'inconclusive' ($n(I)$) information is factorially varied across targets. For novices (i.e., civilian university students), inconclusive information dilutes threat assessments, so that judgments tend to err on the side of "friendly". For experts (i.e., senior Navy officers), inconclusive information accentuates threat assessments, so that judgments tend to err on the side of "foe". Threat assessments are in all cases best fit by a model that includes a component based on the 'balance of evidence' [i.e., $n(H) - n(P)$] and a component based on the scaling of inconclusive evidence [i.e., $n(I)$]. A final study confirms that these judgments are resistant to scenario manipulations that vary the global threat context.

DECISIONS FROM MEMORY: INFORMATION INTEGRATION DEPENDS ON THE INFORMATION FORMAT

Arndt Bröder

Gigerenzer et al. (1999) promote a theory stating that most of our decisions are "fast and frugal" and based on simple noncompensatory heuristics. One of these heuristics is a simple lexicographic rule for inferences based on binary probability cues and is called "Take The Best" (TTB). Within their framework of an "adaptive toolbox" of heuristics, Gigerenzer and colleagues claim that frugal heuristics without information integration are adaptive especially under conditions which require costly information search, for example in memory. We introduce a new paradigm that allows one to determine the presumed decision strategy of participants when attribute information has to be retrieved from memory. Whereas our first two experiments corroborated Gigerenzer et al.'s claim, a third experiment demonstrated that a simple TTB strategy is predominantly employed with verbal cue information, but not with pictorial cue information. Hence, the nature of information integration strategies used appears to depend on the information format.

CONFIDENCE IN THE AGGREGATED OPINION OF CORRELATED JUDGES

David Budescu

This talk focuses on the aggregation process employed by individual Decision-Makers (DMs) when combining probabilistic information from multiple sources. Typical examples are investors who combine forecasts from various financial advisors regarding the chances of certain stocks to appreciate in value, or patients who aggregate information from various experts about the chances of success of a given medical procedure. In all these cases the advisors' opinions are correlated because the forecasters rely on similar (sometimes identical) data, have similar training and share similar values. The talk will review a general algebraic model that assumes that the DM combines information by averaging the various forecasts, and that the DM's confidence in the aggregate is inversely related to the variance of the (possibly weighted) mean forecast. This model is used to derive a series of predictions about the factors that affect and drive the DM's confidence. I will review results of four empirical studies (Budescu & Rantilla, 2000; Budescu et al. 2003) designed to test and validate these predictions.

The experiments seek to characterize (a) the nature of the aggregation rules used by DMs, (b) the factors that affect the DMs' confidence in the final aggregate, (c) the nature and level of dependence between the aggregates and the confidence they inspire, (d) the circumstances under which subjective aggregates are most effective, and (e) the factors that determine the DM's preference for certain advisors.

References:

Budescu, D.V. & Rantilla, A. K. Confidence in aggregation of expert opinions. Acta Psychologica, 2000, 104, 371-398.

Budescu, D.V., Rantilla, A.K., Yu, H., & Karelitz, T.M. The effects of asymmetry among advisors on the aggregation of their opinions. Organizational Behavior and Human Decision Processes, 2003, in press.

JUSTICE PREFERENCES AND THE ARROW PROBLEM

Norman Frohlich & Joe A. Oppenheimer

Arrow showed that there is no general solution to the aggregation of ordinal preferences nor of individual welfares into either a sensible social choice or a welfare measure. The standard solution to this problem has been to consider 'spatial models' that are built on spatially restricted preferences (or welfares). We show that an alternative family of solutions exists: one that may in fact be far more attractive. By replacing self-interest with a simple form of other-regarding preferences, coupled with a sense of justice, we establish the possibility of sensible aggregate choice. Various assumptions regarding the form such considerations may take in the utility function are discussed, and for a number of models, we establish conditions for a Condorcet winner in a problem of pure redistribution: one that has been used to establish the plausibility of cyclic outcomes.

PRINCIPLES AND METHODS FOR COMBINING EXPERT JUDGMENTS

Andrew Gelman, Robert Clemen, Roger Cooke, James Hammitt, David Krantz, and Francis Tuerlinck

Various methods have been proposed for combining expert judgments, including empirically-weighted averages and statistical models such as copulas that treat experts'

numerical estimates as data. We study the properties of some of these methods in some simple example and discuss their relation to normative standards including unbiasedness, calibration of predictive probability statements, and Bayesian analysis of direct data.

LEARNING IN DYNAMIC DECISION MAKING

Cleotilde Gonzalez

Dynamic Decision Making (DDM) involves a series of multiple and interdependent decisions made in real-time in a continuously changing, autonomous environment (Edwards, 1961; Brehmer, 1990). Achieving control of a dynamic system takes lengthy task practice. This process requests the decision maker to have clear causal and temporal models of decisions and outcomes. An instance-based approach to decision making suggest that people learn by accumulating and refining the utility of examples, each containing the decision-making situation, action, and its expected utility at the moment of the decision (Gonzalez, Lerch and Lebiere, forthcoming). In this presentation I will show empirical evidence of instance-based learning in DDM. I will also show the development of a cognitive model of decision-making in DDM that demonstrates some of the learning mechanisms decision makers might be using in dynamic situations. I will conclude with proposed mechanisms to support learning in DDM.

SOURCE EFFECTS IN ADVICE TAKING

Nigel Harvey

I shall review some recent research concerned with how the source to which advice is attributed affects the way that the advice is assessed and used. First, the extent to which advice is assessed as accurate does depend on the quality of the advice. However, the extent to which it is assessed as trustworthy depends interactively on the quality of the advice and on the extent to which the assessors judge that they share values with the source. Second, ability to use advice can depend on the distinctiveness of the different sources: presumably, it is easier to remember which sources of advice are good and which are poor when they are more distinct. Under certain conditions, however, people are better able to use the advice when the good and poor advisors are less distinct. We argue that this occurs when the primary aim of the users of the advice is to ensure that the combine information from diverse sources.

THE ROBUST BEAUTY OF THE MAJORITY RULE

Reid Hastie and Tatsuya Kameda

This talk summarizes recent simulation and behavioral studies of the accuracy of group choices made under eight group decision rules. The comparatively superior performance of the Majority Rule is demonstrated across several variations in Monte Carlo simulated test bed environments. The discussion attempts to explain the relatively excellent performance of the Majority Rule.

INVESTIGATION OF A MODEL OF OPTIMAL INFORMATION INTEGRATION IN GROUPS

Verlin B. Hinsz

A model based on signal detection theory proposes how groups might ideally integrate available information. This ideal group model predicts optimal levels of group performance. Analysis of data reveals that a sample of groups performing a recognition memory task achieved optimal levels of information integration and efficiency as specified by an ideal model. Surprisingly, analysis at the level of each group suggested that half the groups exceeded optimal levels of performance. I hope discussion will help explore ways of accounting for the finding of supra-optimal group responses.

ASYMMETRIC ADAPTABILITY: DYNAMIC TEAM STRUCTURES AS ONE-WAY STREETS

John Hollenbeck

We introduce the concept of "asymmetric adaptability." Using 63 teams that worked on a complex and dynamic information processing task, we replicated the common cross-sectional structural contingency finding that functional structures perform better within predictable environments, whereas divisional structures performed better within unpredictable environments. Unlike most tests of contingency theories, however, we directly tested whether teams could actually adapt in the manner directly implied by the theory and found evidence that one type of change is more difficult than another. Teams responded significantly more favorably to Functional-then-Divisional shifts than they did to Divisional-then-Functional shifts, and these reactions were exacerbated when the team was high in general cognitive ability. We discuss the need to complement the static logic behind many contingency theories with a dynamic logic that explicitly challenges an assumption of symmetrical adaptation.

STRATEGIC INFORMATION SHARING IN GROUP DECISION MAKING

Andrea B. Hollingshead

Laboratory research on cooperative group decision making has largely ignored the role of members' goals on information sharing and group outcomes. In the vast majority of studies, it has been assumed that all group members have the same objective: to choose the best decision alternative for the group. Individual and group incentives are aligned such that members achieve better outcomes when the group makes a high quality decision. The underlying assumptions in this dominant paradigm may be violated in many natural group decision making situations where groups are often composed of members with different goals and incentives. I will present data from a set of studies conducted in collaboration with Gwen Wittenbaum that examine the impact of group members' goals and the underlying incentive structure on information sharing and outcomes in cooperative decision making situations. The data support the idea that information exchange in decision-making groups is a deliberate and strategic process. Group members intentionally determine what information to share and how to share it in order to satisfy goals evoked by features of the decision making context. When aggregating information in group decision making experiments, researchers should take note not only of what information is presented, but also how it is presented and with what effects.

FORECAST NORMS, FORECAST ERRORS, AND THE USE OF METACOGNITIVE JUDGMENT

Dave Krantz

Forecasts arise from mathematical models or from intuition. There are obvious and vast differences between these two sources, but also close linkages between modelling and intuition: models are constructed by people, using natural inductive and deductive processes, and in turn, familiar models can guide and refine intuitive reasoning processes. If the best model that could be constructed, for a given forecast problem, is viewed as a norm for judgment about that problem, then not only does human intuition err, departing from this norm, but sophisticated model construction processes may err likewise, leading to models and forecasts that are worse than the norm. One can classify and perhaps understand many of the shortcomings of intuitive judgment by noting analogous shortcomings that can occur in the construction of mathematical models. Model construction can be analyzed into a number of different segments, and errors are possible in each segment. This analysis produces a classification of errors that is also useful for intuitive forecasts. A model-based forecast is often accompanied by a report that summarizes the process of model construction and validation. Such a report can include multiple measures indicating the level of confidence to be placed on various aspects of the process, and perhaps on the forecast itself. The report may guide decision making based on the forecast, and in particular, when several different forecasts are available, such reports may help to resolve conflicts. Analogously, intuitive forecasts might also be accompanied by detailed process reports, such as think-aloud protocols, or by simpler metacognitive reports. Pushing this analogy farther, one might expect metacognitive reports that refer to several different aspects of the intuitive forecast process. For example, reports of adequate fit for a "mental model" might be distinguished from reports of uncertainty about "mental parameters" or reports of "mental sampling bias." Process or metacognitive reports may be useful in resolving conflicts when different judges offer distinct intuitive forecasts. One of the implications of this idea is that metacognitive judgments should be elicited using careful explanations and instructions, to isolate different aspects of the intuitive forecast process. Dimensional labels such as "probability" and "confidence" may be inadequate instructions for assessing the details of the intuitive forecast process.

LAY INTUITIONS ABOUT COMBINING QUANTITATIVE JUDGMENTS

Richard P. Larrick and Jack B. Soll

Averaging estimates is an effective way of improving accuracy when combining expert judgments, integrating group members' judgments, or using advice to modify personal judgments. If the estimates of two imperfect judges ever fall on either side of the truth, which we term *bracketing*, averaging must outperform the average judge for any weakly-convex penalty function. We hypothesized that people often hold an incorrect theory about averaging, falsely believing that the average of two judges' estimates would be no more accurate than the average judge. A series of experiments demonstrated this misconception. The robustness of this misconception was tested by examining lay intuitions in a collectivistic, holistic culture (China) and by manipulating the salience of bracketing. We conclude by describing how every day life provides few opportunities to learn the benefits of averaging and how misunderstanding averaging contributes to poor intuitive strategies for combining estimates.

IMPLICATIONS OF A NORMATIVE THEORY OF JURY DECISION MAKING

Barbara Mellers

Despite the relatively small number of jury trials, the impact of jury decisions is felt far and wide. Assumptions about what juries would do if faced with the decision are used to justify settlements, fines, and more. Using expected utility theory, Bayes' Theorem, and signal detection theory, we offer a normative account of jury decision making. We then ask what values of d' (discrimination) must jurors have to achieve various error rates. Results suggest that juries are probably making more mistakes than any of us would like to think.

DECISION NETWORKS AND ORGANIZATIONS

Kirshna Pattipati

This talk provides an overview of engineering models of decision networks performing distributed hypothesis testing and command organizations executing specific missions. We begin with the problem of modeling a single decision maker (DM) in binary event detection tasks, and show that the expertise of an individual DM can be characterized by a relative operating characteristic (ROC) curve. Then we consider a distributed version of the event detection (hypothesis testing) problem, wherein multiple distributed DMs cooperate as a team to reach a final decision. Key findings in this case are that the aggregated organizational expertise is operationalized by a *team* ROC curve, and that the jointly optimal decision procedures at each DM are in the form of *coupled* operating points on their individual ROC curves. Using the distributed detection paradigm, we illustrate the impact of task structure on the performance of organizations with different designs. We conclude that the architecture of an organization must be matched correctly to its task structure to achieve superior performance, leading to the concept of *congruence*. We elaborate on this concept in terms of a trade-off between decision performance and internal communication, and develop a method for synthesizing congruent organizational structures. This is followed by a discussion on the need to seek a proper balance among task scheduling, resource allocation, and decision hierarchy, and the development of a methodology for modeling missions and synthesizing the concomitant optimal organizations. Finally, we conclude with a summary of current results in adaptive organizations and future research directions.

AGGREGATION OF PROBABILISTIC ORDINAL PREFERENCES

Michael Regenwetter

Much of social choice theory is dedicated to the elicitation and aggregation of ordinal preferences. Voters may experience uncertainty as to which vote to cast, especially when ballots are complex and when the number of candidates is large. Pollsters and election researchers may experience uncertainty about voter preferences. Furthermore, as the whole world witnessed in 2000 U.S. presidential election, election officials may experience uncertainty as to which ballots actually were cast. It is therefore very natural to investigate social choice theory when ballots are viewed as probabilistic. The present talk will provide an introduction and overview of descriptive probabilistic models of social choice behavior as well as their application to real world election and survey data. A major emphasis will be placed on the new insights that this "behavioral"

approach may yield regarding policy implications and regarding the wise choice of "good" election methods.

GRAPHICAL MODELS AND COLLECTIVE CHOICE

Whitman Richards

Graphical models have led to important advances in probabilistic reasoning because knowledge about relationships in the domain constrains conditional inferences (Pearl, 1988.) In a similar manner, a graphical model can constrain preference orderings of agents engaged in decision-making. As long as all individual rankings of alternatives are consistent with one global graphical model, then certain model forms guarantee a stable outcome that will survive counterproposals. We explore conditions on model forms and on individual rankings that almost always result in stable outcomes and those that do not. (Collaborators: D. Richards & B.D.McKay.)

Intuitive Strategies for Revising Opinions: Are people appropriately influenced by others?

Jacks Soll

In estimation and forecasting, simply averaging judgments is a highly effective way to reduce error. However, in a series of studies we find that people typically do not average. Rather, many people employ the choosing strategy. They make a guess about which person (in our work it is always out of two people) knows more about a given question, and put high weight on that person's judgment. People do occasionally average, and our data suggest that this happens when they are not sure about who is more expert. Because people often choose and occasionally average, the resulting distribution of weights is W-shaped. Intermediate weights that might be expected from an anchoring and adjustment process are relatively rare in our data. Averaging often works well, but sometimes it does not. We propose a model that describes the conditions under which averaging is a good strategy. A key element of the model is the bracketing rate, which we define as the proportion of questions on which the estimates of the two judges fall on either side of the truth. If judges are unbiased and have uncorrelated forecast errors, we would anticipate a bracketing rate of about 50%. In our data, the rates are typically around 40%, which is high enough to exploit the benefits of averaging. The relative effectiveness of averaging also depends on the MAD ratio (which measures the relative accuracy of the two judges) and the probability of detecting the expert. It is interesting that averaging outperforms choosing under a wide range of parameter values. Finally, the model does an excellent job of accounting for our data.

EFFECT OF MAJORITY RULE AND INITIAL BIAS ON INFORMATION AGGREGATION BY GROUPS

R. D. Sorkin, S. Luan, and J. Itzkowitz

We simulated the behavior of juries using a distributed-detection model of group information aggregation and deliberation. The model assumes that each juror observes the evidence with a unique level of expertise and an initial bias. Upon hearing the yes/no votes of the other jurors, each juror revises her own vote in a Bayesian fashion. The voting process iterates until the votes for or against conviction reach a threshold majority (simple, two-thirds, three-fourths, or unanimous majority) or time runs out. The vote

updating process assumes juror knowledge of the level of expertise and bias of the other jurors. We also evaluated the performance of juries who used non-optimal updating rules, such as Delphi and simple sub-majority-shift (social decision scheme) rules. The resulting performance (accuracy and bias) of our simulated juries was examined for different levels of initial (average) juror bias. Results indicated that jury performance depended on the jury's size, required majority, and initial bias. The most accurate group performance and smallest decision bias were produced by the unanimous majority rule. The number of correct decisions increased as the required majority (simple, two-thirds, three-quarters, etc.) was made more stringent. Juries that began deliberation with a neutral bias, used a Bayes' updating rule, and required a unanimous majority, produced detection indices very close to the theoretical optimum.

GROUP POLARIZATION AND COLLECTIVE INFORMATION POOLING: DOES POOLING PROMOTE POLARIZATION?

Garold Stasser

A pervasive finding in the collective judgment literature is that group discussion accentuates popular tendencies. That is, group judgments and individual judgments after discussion are more polarized than are individual judgments before discussion. One class of explanations for this phenomenon rests on the assertion that discussion content drives polarization. Most notably, Vinokur and Burstein's (1977) Persuasive Arguments Theory (PAT) uses an information sampling metaphor and ascribes a critical role to *novel* arguments, arguments that relatively few people know. Core ideas in PAT are: 1) discussion allows for people to exchange arguments; 2) these arguments are sampled from the same pool that gave rise to pre-discussion sentiments and thus will tend to support these initial sentiments; and 3) discussants will bring some novel arguments to the table and these novel arguments are particularly persuasive and promote more extreme judgments. The collective information pooling literature documents another kind of pervasive finding that seems to contradict the idea that novel or unique information is instrumental in shaping collective responses. For example, group discussions en route to a collective choice are dominated by widely-shared information and unique information is often omitted from consideration. Thus, group decisions typically reflect the common knowledge of group members and not their unique knowledge (Gigone and Hastie, 1994). The Collective Information Sampling (CIS) model (Stasser and Titus, 1977), like PAT, uses a sampling metaphor to understand why groups seemingly overlook or ignore unique information. Explorations using DISCUSS, a computational model of collective information sampling and choice (Stasser, 2000), illuminate the similarities and differences between PAT and CIS and the phenomena they purportedly represent. For example, one important difference is the representation of how information access is distributed among group members. In PAT, an argument is novel because people are not likely to think of it although all group members have the same likelihood, albeit low, of contributing the novel argument. In CIS, information is unique because only one or a few members have access to it before discussion. Thus, most of the group members cannot contribute a unique item to discussion. Moreover, attempts to model polarization using the sampling/resampling processes suggested by PAT are not sufficient to generate polarization. Plausible auxiliary components to PAT are considered.

AVERAGING PROBABILITY JUDGMENTS

Thomas S. Wallsten and David Budescu

The central tendencies of N probability judgments can be more diagnostic or better calibrated than any of the individual judgments, depending on the extent of their inter-correlation. We review two theorems, as well as extensive computational and empirical research, that clarify and establish this result.

BAYESIAN INFORMATION AGGREGATION: MODELING ISSUES AND CHALLENGES

Robert Winkler

Bayes' Theorem provides a natural framework for information aggregation in decision making. In principle, it can handle any sort of information aggregation in a logical, sensible manner and can support rational decision making in a normative sense. However, in practice many interesting and challenging modeling issues can arise in applying the Bayesian approach. In my talk I will focus on some of those issues, with particular attention to the aggregation of probabilities.

RECEIVING ADVICE AND MAKING JUDGMENTS IN A SOCIAL ENVIRONMENT: A PRAGMATIC APPROACH

Ilan Yaniv

Using advice is a basic practice in making real life decisions, although, until recently, relatively little attention has been given to it in either empirical studies or theories of decision making. Decision makers integrate the opinions of several other individuals. Two theoretical perspectives are invoked. The first perspective is focused on the social environment including, (i) the giver of the information (e.g., expert, advisor, forecaster), and (ii) the receiver (decision maker, judge). The interaction between givers and receivers is analyzed in strategic and pragmatic terms (Grice's rules of conversation: informativeness, relevance, truthfulness). A second theoretical framework is introduced which is based on an analogy between judgments and attitudes. Insights from the study of attitude change (communicator-audience) are used to explain the influence of advice. In a series of studies we investigated people's weighting policies for advice. Respondents were asked to provide final judgments on the basis of their initial opinions and advice presented to them. The respondents' weighting policies were inferred. Analysis of the advice-weighting policies showed evidence for advice discounting, reputation effects, knowledge effects, and advice distance effects. In addition the consequences of these policies for accuracy are evaluated.

CAST: AGENT-BASED PROACTIVE INFORMATION EXCHANGES FOR TEAM DECISION-MAKING

John Yen

Psychological studies about human teamwork have suggested that effective human team can often anticipate needs of teammates and proactively offer relevant information to them based on a mental model shared by the team. In this talk, I will describe an ongoing MURI research effort to develop an agent architecture (called CAST) that enables agents in a heterogeneous team (which could include robots,

software agents, and people) to anticipate information needs of other teammates and to proactively exchange information for addressing their needs. These capabilities are enabled by capturing knowledge about the structure and the process of a team using a high-level language (called MALLEt). This teamwork knowledge is compiled into a Predicate Transition Net, which serves a key component of the agent's shared mental model. CAST achieves high efficiency of teamwork without sacrificing its flexibility. The theoretical foundation of CAST has been established based on two formal agent theories: SharedPlan theory and Joint Intention theory. This foundation has formally specified the semantics of proactive communications, and offers opportunities for extending agent communication protocols for supporting proactive teamwork. Preliminary experimental results suggest that CAST is a promising approach for modeling proactive information exchanges using agents. Based on CAST, we are currently developing an agent-based training system for an AWACS-like synthetic task (i.e., DDD).